INVITED PAPER

Energy Challenges in Current and Future Optical Transmission Networks

The authors believe that future communication networks should be efficient, affordable, and sustainable, and that the information and communication sector can help to reduce energy consumption in other sectors.

By Daniel Kilper, *Senior Member IEEE*, K. Guan, K. Hinton, and R. Ayre

ABSTRACT | In this paper, we examine how energy constraints might shape future optical communication networks and the impact that current technology trends may have on future energy use. Historical factors and prevailing complications associated with fiber capacity point to an increased focus on energy to enable tighter photonic and electronic component integration and larger networks. Energy requirements and associated challenges are described at the component, system, and network level.

KEYWORDS | Energy consumption; energy efficiency; networks; optical transmission

I. INTRODUCTION

Energy has become a topic of increasing importance for optical communication systems in recent years. This trend has been fueled by a combination of factors that include beneficial impacts as well as complications and technological challenges. A 10% increase in broadband penetration has been shown to raise the per capita economic growth of a nation by 0.9%–1.5% [1]. Thus, ubiquitous access to broadband will have a major impact on economic growth which, in turn, may potentially incur a corresponding increase in a nation's carbon footprint [2]. Added to this will be the power consumption of the broadband equipment itself, including the home and network equipment. The role of communication networks on global energy consumption must be taken seriously. The growth that is concomitant with communication technologies is bringing focus to both scalability and sustainability for optical networks.

Information and communication technologies (ICT) in general can be an important technology to address global climate change and realize sustainable living. Smart technologies use ICT to better monitor and control carbon emissions in everything from home appliances to the electric power grid. Many applications that use networks, particularly when considered from a full "cradle-to-grave" lifecycle analysis, show reductions in carbon footprint even with an increase in the carbon footprint contributions of the network.

One such study suggested that by 2020 the continued expansion of the use of ICT, including networks, has the potential to offset carbon emissions in other sectors of the economy by an amount that is five times ICT's own total carbon footprint [3]. As a result, the increasing focus on climate change and sustainability has brought attention to networks as an enabling technology to address these issues. Furthermore, optical communication systems in many applications, including access and core networks, have been shown to be the most energy-efficient communication platform [4].

Realizing these benefits from optical networks relies on the continuation of scalable, energy-efficient network growth. Over the past decade, data networks have grown in the numbers of users, hosts, and addresses and shown near-exponential traffic growth that is expected to

Manuscript received August 12, 2011; revised October 26, 2011; accepted November 7, 2011. Date of publication March 6, 2012; date of current version April 18, 2012. This work was conducted as part of the GreenTouch Consortium.
D. Kilper and **K. Guan** are with Bell Labs, Alcatel-Lucent, Holmdel, NJ 07733 USA (e-mail: dan.kilper@alcatel-lucent.com; kyle.guan@alcatel-lucent.com).
K. Hinton and **R. Ayre** are with the Center for Energy Efficient Telecommunications at the University of Melbourne, Melbourne, Vic. 3010, Australia (e-mail: k.hinton@unimelb.au; rayre@unimelb.au).

Digital Object Identifier: 10.1109/JPROC.2012.2186102

continue even in mature markets at 30%–50% per annum over the next decade [5]. Network energy efficiency must improve at similar rates in order to avert a rapid increase in the network energy consumption. In 2007, ICT was estimated to make up 1%–2% of the worldwide carbon footprint, and network infrastructure made up roughly 25% of that contribution [3]. While this is proportionally small, it is still comparable to industries such as the airlines. In some countries, telecommunications is among the top three most energy consuming industries [6]. Furthermore, other more energy-intensive industries globally are not growing as rapidly as data networks. Assuming no improvement in equipment efficiency, the energy consumption of communication equipment following current trends would dominate energy use in just one decade [7].

Historically, data networks have scaled through capacity gains from optical transmission systems and Moore's law scaling in the network electronic infrastructure such as Ethernet switches and Internet protocol (IP) routers. Although Moore's law scaling is expected to continue for another decade or longer, the associated energy scaling is already considered to have stalled, creating a thermal density bottleneck that in the computer industry has driven the move to multicore processors. Recent studies of telecommunications equipment indicate that efficiency improvements measured in terms of the energy per bit transmitted or processed are slowing [5]. Historical trends of 15%–20% per annum may fall to 10% per annum or lower [8]. The difference between traffic growth and equipment efficiency improvements has already resulted in thermal density bottlenecks in central office equipment. High-capacity communication equipment, from optical transmission hardware to core routers, is at or above recommended thermal density limits. Therefore, future capacity growth is contingent on maintaining the current per rack thermal loads. Increasingly, higher capacity solutions will likely involve multirack solutions, covering a larger spatial footprint and higher energy consumption. Without aggressive efficiency improvement, the carbon footprint for the network has the potential for dramatic growth.

Efficiency measures applied to the network equipment are also desirable because they can deliver compounded energy savings through reducing thermal management requirements, possibly even enabling passive cooling, and reducing load proportional losses in the electrical power delivery. Telecom centers are generally placed close to the customers they serve, minimizing the length of access cable runs as well as network latency. Thus, it may not be possible to move telecom centers close to the power generation source, and significant power can be lost in the electrical power distribution network. Furthermore, improving the efficiency of the equipment can enable wider use of renewable energy sources, such as solar or wind, in situations that might otherwise not be possible.

Energy has always been central in fundamental considerations of optical communications, usually expressed in terms of the number of photons per bit at a given wavelength. Furthermore, energy consumption has been studied in power-constrained applications such as submarine and deep-space communications. In terrestrial systems, however, energy consumption has been of secondary interest and energy efficiency has usually come as a byproduct of increased bandwidth or other network scaling considerations. This has changed in recent years. Fig. 1 shows the results of an IEEE Xplore database metadata query (title and abstract) on the phrases "energy consumption," "power consumption," "energy efficient," or "energy use" AND'ed with the corresponding search criteria in the legend: "optical network," "optical transmission," and "wdm." In each case, there have been few publications prior to 2007. Following 2007, the number of publications increased rapidly. Publications on the subject of optical networks shows the largest increase, which may be expected due to the broader scope, but also may be reflective of research interest. Much of the work prior to 2007 was on particular devices or technologies that exhibited low-power or efficient operation, in most cases identified as a secondary benefit. Leading up to 2007, there is increasing interest in optical interconnects as a potential solution for mitigating the power density challenges in high-speed electronic circuits or backplanes. During 2007, several papers were published on the overall energy consumption of the network infrastructure in the Internet, including routers, switches, and transmission systems, marking the start of the current research trend.

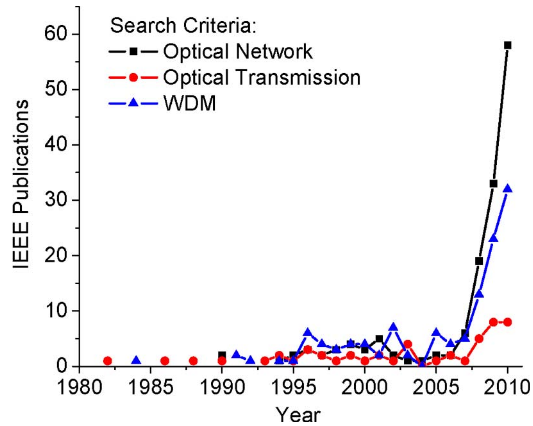

Fig. 1. ***IEEE publications with metadata containing either "energy consumption," "power consumption," "energy efficient," or "energy use" AND'ed with the search criteria shown in the legend.***

Several reports from prominent environmental groups came out starting in 2008, including the Climate Group sponsored by the Global e-Sustainability Initiative [3], World Wildlife Federation in collaboration with the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) [9], and Greenpeace International [10], stimulating further interest from the opposite perspective of using the network to support sustainability measures throughout society. Related activities have been identified as "greening by ICT," whereas efforts to reduce the carbon footprint of ICT equipment is referred to as "greening of ICT" [11]. Research on greening of ICT was further stimulated through inclusion in European Framework 7 calls for proposals along with funding initiatives in Korea and Japan. The greening of ICT networks has become a focal point for industrial and academic research consortia and centers including Intelligent Energy Aware Networks (INTERNET), Center for Integrated Access Networks (CIAN), Mobile VCE, Energy Aware Radio and Network Technologies (EARTH), and GreenTouch. In 2011, the World Resources Institute sponsored the Carbon Trust to develop a set of ICT industry guidelines as part of the global greenhouse gas emissions protocol [12].

In this paper, we examine the issues related to energy use in optical networks that are relevant today and that are anticipated to be important for future networks. In particular, we focus on optical transmission networks, which we define here as terrestrial wavelength-division-multiplexed (WDM) transmission-system-based networks common in metropolitan, regional, and long-haul applications. Other types of optical networks, including passive optical networks, data center networks, and mobile backhaul networks, are not considered in detail, but may share many of the same issues. These access- and enterprise-based networks tend to focus more on aggregation than transmission and therefore do have unique attributes. In this work, we start with a review of high-level attempts to quantify energy use in ICT overall and consider the energy-efficiency metrics relevant to different optical network technologies. We provide an overview of the key energy related issues shaping optical transmission networks. The fundamental limits to energy use in optical transmission are briefly reviewed. A sample commercial off-the-shelf system is constructed and broken down in terms of energy use to the line card and subsystem level. This analysis is followed by an overview of the energy-dependent factors important to line system transmission design and the relationship between energy efficiency and spectral efficiency. At the network level, optical switching is considered along with the enabling effect of optical transmission in facilitating energy savings in the higher network layers. Finally, we review the service- and application-dependent energy issues, which are expected to increasingly impact optical systems.

A. Metrics and Measures

A 1999 paper estimated the power consumption of personal computers (PCs) and network equipment as 8% of the United States' electricity consumption [13], and forecast that, by 2010, this figure would grow to about 50%. This estimate was strongly contested and several subsequent papers have reported substantially lower values, ranging from about 0.5% to 3% (see Table 1). Until the paper published by Baliga *et al.* in 2009 [7], models for power consumption of ICT used estimates of deployed equipment volumes based upon equipment inventory or sales figures. Taking a different approach, Baliga *et al.* applied general network design rules to establish the equipment use, and consequentially power consumption, of a minimal network that would be required to accommodate all the traffic arising from a population of end users, each with a given average demand for access capacity. Subsequently, a "transaction-based model" was introduced which estimates the power consumption of a range of widely used network services (including: general web services, peer to peer, video, mobile voice, and mobile data services), and used energy trends to project the results to 2020 using technology-evolution-based models [14].

Table 1 Selection of Reports on Power Consumption of ICT Over Recent Years. The Range of Values for "% Regional Electricity Use" Arises From Variation on the Data, Method of Calculation, and Range of Equipment Included in the Estimate

Authors	Year	% regional electricity use	Region	Ref.
Huber & Mills	1999	13%	USA	[13][b, c]
Koomey et al.	1999	2%	USA	[15][b, c]
Kawatomo et al.	2001	3%	USA	[16][b, c]
Turk	2001	0.5 - 1.7%	Germany	[17][b, c]
Barthel et al.	2001	0.9 – 1.5%	Germany	[18][a, b, c]
Roth et al.	2002	< 2.3%	USA	[19][a, b, c]
Cremer et al.	2003	7.1%	Germany	[20][a, b, c]
Baliga et al.	2009	0.4%	OECD	[7][d]
Lange et al.	2010	Not estimated	Global	[21][a, d]
Vereecken et al.	2010	Not estimated	Not given	[22][d]
Kilper et al.	2011	Not estimated	USA	[5][a, e]

[a]Includes wireless access
[b]Includes PCs, office equipment, and servers
[c]Users & equipment estimate
[d]Network design model
[e]Transaction based model

Although the percentages in Table 1 do not appear to be high, we need to keep two factors in mind. First, the percentages are growing at a rate commensurate with the growth of Internet traffic, which is estimated to be approximately 40% per annum. Therefore, without improvements in technology, it will not be long before these percentages become unsustainable. Second, the power consumption represented by the values in Table 1 is concentrated in a relatively small number of facilities. This means these facilities must be provided with a large electrical supply as well as dissipate a significant amount of heat. As will be discussed below, the issue of supply and heat dissipation has already become a major challenge in the ICT industry, including optical transmission systems.

Due to the high growth rates across many aspects of network evolution, energy efficiency is a central metric for most greening of ICT activities. The definition of energy efficiency, however, is dependent on the particular technology or application. The energy per bit metric is widely used for both transmission and computation elements within a network. For transmission, this is often written as the power per bit per second. At the device level, defining a bit or data rate is usually straightforward. At the system or network level, the definition is complicated due to the presence of overhead data and dropped packets. For systems, using the information bitrate removes the coding overhead. Likewise for networks, the considering network "goodput" traffic restricts the bits to the data delivered to the end user or application. Separating the goodput from the network traffic, however, may not be straightforward as the equipment that is reporting the traffic volume often cannot make the separation of overhead in different network layers.

Any improvement measure that focuses on efficiency raises the concern of rebound effects characterized by the Kazoom–Brookes postulate or Jevon's paradox [23], [24]. This phenomenon occurs when energy-efficiency improvements in a particular technology or device result in either direct increases in use or indirect growth elsewhere that lead to a global increase in energy consumption. The positive impact of broadband on economic growth, if it leads to a global increase in energy consumption, would be an example of an indirect rebound effect. Addressing rebound effects, however, is often a matter of policy and pricing, which are largely out of the domain of technical research. Indeed, efficiency measures are an important component of most climate change and sustainability strategies, but technology improvements need to be coupled with policy and investment measures as well [10], [12]. Aggressive use of ICT to offset consumption in other sectors—greening by ICT—is a potential answer to indirect rebound effects. To the extent that new technology can aid us in monitoring the use of the network itself, there is likewise the possibility for technology innovation to address direct rebound phenomena by helping to moderate traffic and network growth as the efficiency improves. However, it is not clear to what extent energy is influencing network growth today, and if it is not a major factor, then these rebound phenomena do not apply.

There are a number of ways in which network equipment energy consumption is reported, and care may be needed in making comparisons. Most network equipment is bidirectional and therefore includes both transmit and receive capabilities. Sometimes, however, the equipment efficiency is reported using unidirectional data rates, resulting in half the reported energy per bit. It is also important to consider the time dependence of both the power (or energy) and traffic. Both power and data rates are usually reported using maximum or peak values. However, the mean values are more reflective of actual use.

Energy in networks can also be studied from a lifecycle perspective. A cradle-to-grave approach is often used when considering greening by ICT or ICT-enabled smart applications. A cradle-to-grave analysis starts with the raw materials used in the telecommunications equipment, and includes the equipment manufacturing, installation, maintenance, use over the network life, final decommissioning, and disposal. This analysis can be focused around a particular service that uses the network, such as distance learning or video conferencing. One of the main challenges with such a service-oriented analysis is to accurately allocate the network carbon intensity associated with the service under investigation. Networks support many different services and the path that data take through the network can cover multiple continents and vary widely over time. Often the core network, which carries the greatest uncertainty in this respect, contributes negligibly to the overall service footprint. Furthermore, for high-capacity telecom equipment, the embodied energy is typically around 10% of the total and thus most attention is directed to the use phase [12]. Exceptions include cases such as greenfield access networks that require trenching or other energy-intensive installation [25].

II. KEY ISSUES

The key issues for the networking industry are thus the challenges of scaling networks to meet the demands for increasing capacity, while minimizing the growth in network energy consumption. Continued network growth is expected to have dramatic consequences on the design of optical communication networks and their associated technologies. Optical channels are rapidly approaching fiber capacity limits. A recent analysis indicated that laboratory systems are within a factor of two of capacity limits [26]. Continued capacity expansion will therefore require the use of more bandwidth. This bandwidth can come from expanding the transmission band within a single fiber, using additional degrees of freedom such as polarization and phase, or moving to additional fibers or transmission modes within a fiber. Regardless of the approach, this shift is essentially a move from a single line system architecture to a multiple parallel line system architecture—similar to the move to multicore architectures in electronic processors. The scalability of such systems will be largely driven by size (footprint) and energy, with the two also being coupled. A 3-cm-diameter fiber cable can house 1000 fibers. Assuming progress in the development of technologies such as multicore fibers or multiple spatial mode multiplexing each fiber can potentially support 10–100 cores and/or spatial modes, as well as several optical amplification bands. This could provide as much as five orders of magnitude of intrinsic capacity growth. The challenge, therefore, is to realize amplification and transceiver technologies that can exploit this bandwidth without causing a five-orders-of-magnitude increase in office/circuit pack footprint or energy. Photonic and electronic integration will be central to meeting this challenge.

A 10-Gb/s transmission line card today consumes of the order of 50 W, but has a wall plug to output optical signal power efficiency of less than 10^{-3}. In contrast, recent analysis of a potential minimum-energy hardware implementation suggests the possibility of a 250-mW, 100-Gb/s transceiver implementation using 2020 generation components [27]. Assuming unity spectral efficiency, a full C-band set of 45 wavelengths would consume 11 W. Taking similar performance on the client side, 22 W is still less than half the power of a 10-Gb/s line card today. Implementing both solutions on a 50-W line card would represent a 450 times efficiency improvement. While this analysis is very hypothetical, it points to potential for efficiency improvements that approach 10^3. Thermal density may be a limiting factor here as well as the radio-frequency (RF) interference on a line card carrying 90 sources (client and line side) operating at 100 Gb/s. While these present major research challenges, they are precisely the same challenges that are facing the photonic interconnect community.

For transmission systems, however, the line side performance requires signals that can survive long distance fiber propagation. If signal regeneration is needed along the transmission path, this will multiply the number and consequently power of the transceivers in the system by one plus the number of regenerations. Since the reach of a transmission system is also strongly dependent on the modulation format and channel coding, there is considerable opportunity for design tradeoffs in the transceiver and line system to achieve the highest overall system efficiency for a given reach requirement.

The line amplifiers deployed along an optical transmission link also offer scope for efficiency gains. In the line amplifiers, pump lasers may be integrated and share thermal controls and other overhead. Assuming each pump laser requires 1 W of electrical drive power, then the amplifiers at a single repeater site for a 1000-fiber cable would require 1 kW for the lasers alone. Allowing a factor two for thermal management and other overheads, the resulting 2-kW solution would be 25 times more efficient than current amplifiers on a per bit basis, which use roughly 50 W per fiber. Using the 100-Gb/s solution described above, another factor of five would be realized from the increased spectral efficiency on the line resulting in an overall improvement of 125. The potential is not as large as for the transceivers, but the amplifiers are already roughly ten times more efficient that the transceivers in long-haul systems.

Optical networks will spread geographically as optical access becomes the dominant access technology. It is also likely they will flatten as the same optical equipment finds use in different network segments (e.g., merging metropolitan and long-haul networks). Further, the growing diversity of services offered to customers will mean traffic may vary substantially in time and location. Collectively, these trends may drive the use of more dynamic physical infrastructure to gain efficiency in network utilization, which in turn can translate into energy savings. Dynamically powering equipment up and down based on utilization can enable higher instantaneous capacity for a given mean energy use across a network. This situation is analogous to the use of dynamically powered small cells in a wireless network to enable high mobile data capacities. Often referred to as cognitive networking in cellular systems, efficiency might be gained through greater intelligence to automatically respond to varying traffic conditions through dynamic spectrum and resource utilization. In cellular systems, this includes adapting frequency, modulation format, data rate, transmit power, and antenna patterns. For a wireline optical system such capability might include adapting lightpath, fiber, wavelength, modulation format, data rate, transmit power, and impairment compensation. Each of these factors can influence the network efficiency.

A dynamic network capability in the physical layer also has potential to impact the efficiency of the higher network layers. In general, electronic devices are more efficient for high-speed switching and processing than other technologies. However, for high data rate signals, if the number of switching operations per bit is small or the switching is sufficiently slow, then optical switching can be more energy efficient. Circuit switching, which involves occasional rerouting of connections or time multiplexing within a wavelength, has been proposed as an energy-efficient architecture [28]. Commercial routers and electronic cross connects widely used today perform operations that include complex address processing, packet pattern matching, security processing, data buffering, and protocol processing operations that are not efficiently implemented using optical techniques. These operations are high power consuming, and therefore, any opportunity to limit unnecessary processing, scale the electronic clock speeds, or use dynamic power functions such as sleep modes can lead to large network efficiency gains—even in cases in which any consequential transmission-layer optical system power consumption increases. Dynamic operation of electronic switching and processing can be implemented without using dynamic functionality in the lower network-layer optical systems. However, leaving the optical systems fully powered for maximum capacity while the client hardware has been scaled to lower rates or turned down is an inefficient use of optical systems. Furthermore, a dynamic physical layer could enable resource pooling to achieve higher instantaneous capacity and be used to achieve yet greater efficiency in the higher layers. Recent studies have examined the use of high-capacity optical connections to gain greater efficiency in data center management and content placement [29], [30]. For example, dynamic wavelength capabilities have been considered for use in moving content between data centers to take best advantage of renewable energy sources, reconfiguring the network to rapidly transfer large amounts of data [31].

III. FUNDAMENTALS OF TRANSMISSION ENERGY

For reliable transmission in any communication system, the receiver must be able to distinguish the desired signal from the inherent noise in the system, which may include an assortment of interfering signals. The received signal power must be high enough to enable the message to be recovered despite the noise power, using some form of signal processing if necessary.

Shannon's theorem describes the relationship between the achievable capacity of a transmission system, the system bandwidth, and the ratio of the signal power to the noise power at the receiver required for error-free communication. This relationship is derived from thermodynamic principles and is set by the minimum change in entropy for a dissipative system [32]. The theorem sets a fundamental lower limit to the energy consumption of a dissipative or nonadiabatic transmission system, i.e., systems in use today. However, in practice, the energy required to prepare the signal for transmission and recover the signal far exceeds this inherent limit. Furthermore, the nature of the channel will be different depending on the characteristics of the modulation and noise. For example, coherent modulation in a system that uses phase-insensitive amplification, as is typical in commercial implementations, will form an additive white Gaussian noise (AWGN) channel. On–off keyed amplitude modulation however forms a channel with non-central-negative binomial distributed noise [33]. An ideal photon channel or photon number state exhibits quantum limited performance that approaches an AWGN channel in the limit of a large number of photons. The ideal photon channel has the unique property that the spectral efficiency continuously decreases for decreasing signal-to-noise ratio (SNR) with no lower bound [34]. In general, for a given channel, the minimum received SNR can be used to determine the minimum signal power or energy per bit.

Physical limits for a channel with phase-insensitive amplification have been examined considering a range of different system and technology constraints applicable to fiber optic transmission systems. The noise properties and minimum receiver sensitivity were derived for a system with multiple erbium-doped fiber amplifiers (EDFAs) [33]. While distributed amplifiers minimize the amplified spontaneous emission (ASE) noise, EDFAs are used in discrete configurations that for large gain are limited to a minimum noise figure of 3 dB [34]. For this case, the minimum energy per bit of the transmitted signal depends on the square of the number of amplifiers [27]. The overall energy consumption per bit is minimized when the optical line amplifier spacing is set so that the total power consumed by all the amplifiers equals the power consumed by the transmitter and the receiver for that link.

IV. ENERGY IN OPTICAL TRANSMISSION SYSTEMS

A common metric for optical transmission systems is the cost per bit per second per kilometer, which emphasizes the importance of capacity and reach. The cost constraint is usually connected with resource utilization, such as the number of transponders or amplifiers. Improved energy efficiency typically follows from maximizing this metric since the system energy is often proportional to the number of such resources. While cost is clearly the main driver for commercial systems, it is hard to predict or control due to the market influence. Using power instead of cost is also complicated, because the power requirements for commercial systems have historically been dictated by the implementation of specific overhead components such as board controllers and not directly from the power of the optical systems or the optical components.

A. Line System Anatomy

Line systems can be made up of multiple shelves, which hold line cards that perform different functions, including: optical multiplexers/demultiplexers, optical transceivers, optical amplifiers, and wavelength-selective switches. A line system shelf will include a power and cooling system, a shelf control system, and several line cards. Considering a 1-kW shelf, 20% of the power might be lost to shelf level power conditioning and cooling. Shelf controls that talk to network management software and other shelf line cards through the back plane account for roughly 10% of the power. The remainder goes into the line cards themselves. Thus, we can write the total shelf power P_{shelf} as the sum of these contributions

$$P_{\mathrm{shelf}} = P_{\mathrm{F}} + P_{\mathrm{SC}} + P_{\mathrm{L1}} + \ldots + P_{\mathrm{Ln}} \qquad (1)$$

where P_{F} is the fan cooling and power conditioning 200 W, P_{SC} accounts for shelf controller electronics 100 W, and P_{L1} to P_{Ln} are the line card powers, typically 50 W per slot. A transceiver line card will emit roughly 1–10 mW of optical power for each signal and an amplifier line card will emit up to 200 mW. Thus, the overall wall plug efficiency from electrical power in to optical power out is in the range: 10^{-5} to 10^{-3}. The power of each line card P_{L} can be similarly broken down into components, grouped into the common equipment and line hardware, respectively

$$P_{\mathrm{L}} = \frac{1}{\eta}\left[(P_{\mathrm{BC}} + P_{\mathrm{CM}}) + (P_{\mathrm{CC}} + P_{\mathrm{LC}})\right]. \qquad (2)$$

Power conditioning in the line cards η can be 90% efficient. A line card will have a board controller that manages the modules/components and other associated electronics such as backplane line drivers all with power P_{BC} (13 W in our example). Transceiver cards will also include client side optical modules with power P_{CM} (typically 1 W) including both transmit and receive components. These can be short-reach pluggable modules such as small form-factor pluggable (SFP) or multisource agreement (MSA) units. The remaining power is divided into a conditioning component quantity P_{CC} (30 W) and a line component quantity P_{LC}, which includes power that is proportional to the output optical power (1 W). The conditioning portion includes laser temperature control, high-speed electronic processing (clock and data recovery, forward error correction, and potentially optical conditioning such as electronic dispersion compensation). The transmitter laser output, receiver input, and for amplifier line cards, the pump laser are all examples of contributing devices to the line component power.

In the past, energy-efficiency improvements for an optical transmission system have primarily come in two ways: increasing the density of line cards or increasing the line card capacity. Before 2000, commercial line systems had a maximum capacity of 16 wavelengths at 2.5 Gb/s with a total of 40 Gb/s in the fiber and a maximum reach of 640 km over eight spans [35]. A typical shelf was less than 500 W. By comparison in 2011, commercial systems support up to 36 line cards per shelf and 88 channels at 100 Gb/s over distances up to 1600 km [36]. Efficiency improvement was realized by increasing the capacity per line card from 2.5 to 100 Gb/s. Further efficiency gains come at the system level by decreasing the line card footprint to achieve more line cards per shelf. This higher density results in a reduction in the overhead contribution at both the shelf level and the line card level. Considering transmission, the 2011 hardware attains greater efficiency by supporting more wavelengths per amplifier, higher density of amplifier/transmission line cards per shelf, and longer reach.

1) Transceiver Line Cards: The line card conditioning component power P_{CC} can be further broken down into subcomponents for the transceivers and the amplifiers. An example of a transceiver model using external optical modulation is shown in Fig. 2. The transmitter and the receiver will often be housed in a separate module, indicated by the dashed lines, that plugs into the board. For the transceivers, the conditioning components are separated into the laser conditioning, such as temperature control and tuning electronics P_{LaC}, the optical conditioning, such as electronic dispersion compensation P_{OC}, the board electronics including forward error correction P_{EB}, module transmitter electronics P_{ETx}, and module receiver electronics P_{ERx}. Referring to Fig. 2, OPT RX and OPT TX are the receive side and transmit side optical conditioning elements. Note that sometimes these are electronic devices that are included in the Tx/Rx module. These elements can also have a line power component depending on the nature of the technology, such as an optical preamplifier. The total transceiver line card conditioning component power is

$$P_{\mathrm{CC}} = P_{\mathrm{LaC}} + P_{\mathrm{OC}} + P_{\mathrm{EB}} + P_{\mathrm{ETx}} + P_{\mathrm{ERx}}. \tag{3}$$

Referring again to Fig. 2, the transmitter and receiver electronics include the electronic multiplexer (MUX) and demultiplexer (DEMUX) (or serializer/deserializers), and clock and data recovery (CDR), receive side amplifiers (AMP), and modulator drivers (DRV). The transceiver line components include the source laser and modulator P_{SRC} and receiver P_{Rx}. The laser source power is related to the output optical power per bits per second (equivalently the optical energy per bit) through the efficiency η_{L} and the receiver power is related to the input optical power through the efficiency η_{Rx}. The exact form of these efficiency factors will depend on the details of the modulation format. For an externally modulated on–off keyed transmitter the instantaneous optical output power is related to the continuous-wave (CW) laser power P_{CW} times the modulator transmission η_{mod}. The modulator

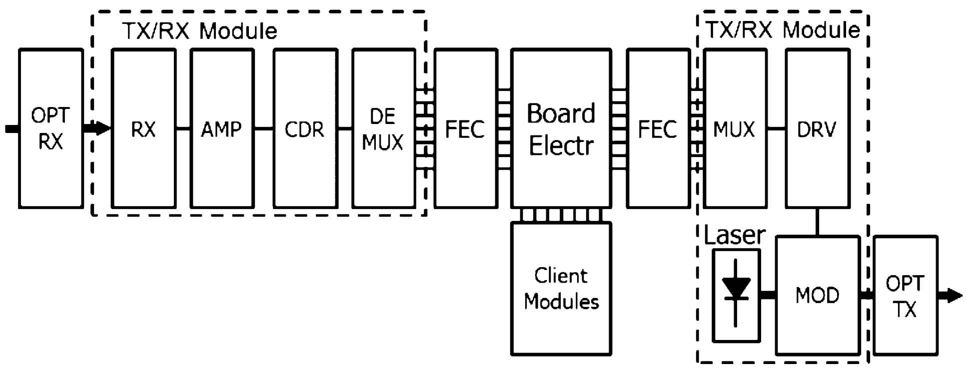

Fig. 2. ***Transceiver line card model for on–off keyed modulation.***

transmission can be related to the drive voltage. For example, with on–off keyed modulation using an external Mach Zehnder modulator, the output optical power is related to the laser source optical CW power, optical losses η_L, and the modulator losses $\eta_{\rm mod}$: $P_{\rm Oout} = \eta_L\ \eta_{\rm mod}$ $P_{\rm CW}$ [27].

Coherent phase-modulated signals might use an in-phase and quadrature (I/Q) modulator on the transmit side and a balanced detector including a local oscillator and polarization diversity on the receive side. Including the source efficiency $\eta_{\rm SRC}$ and receiver efficiencies $\eta_{\rm Rx}$, the full line component power (i.e., related to the transmission signal energy per bit) can be written

$$\begin{aligned} P_{\rm LC} &= \frac{1}{\eta_{\rm SRC}} P_{\rm CW} + P_{\rm Rx} \\ &= \frac{1}{\eta_{\rm SRC}\eta_{\rm mod}\eta_L} P_{\rm Oout} + \frac{1}{\eta_{\rm Rx}} P_{\rm Oin}. \end{aligned} \tag{4}$$

2) Amplifier Line Cards: The amplifier line cards do not include client side optics, therefore $P_{\rm CM} = 0$. The amplifier conditioning components (CCa) fall into the categories of thermal control $P_{\rm TC}$, transient control $P_{\rm Tr}$, and pump control $P_{\rm PC}$

$$P_{\rm CCa} = P_{\rm TC} + P_{\rm TrC} + P_{\rm PC}. \tag{5}$$

On the line side for the amplifiers (amplifier line conditioning LCa), the optical signal power gain is related to the pump power through an efficiency factor η_P and is multiplied by the number of wavelengths k

$$P_{\rm LCa} = P_P = \frac{k}{\eta_P}(P_{\rm Oout} - P_{\rm Oin}). \tag{6}$$

3) Other System Elements: A transmission system may also include other components that contribute to the energy consumption of the system. Both electronic performance monitoring and optical performance monitoring are used for various network management functions. Electronic performance monitoring is usually implemented in a board controller and uses the forward error correction statistics such as frame error counts. Optical performance monitoring includes channel monitors and optical power monitors. Advanced optical performance monitoring to obtain detailed information about the optical signal quality such as group velocity dispersion, optical SNR, or polarization mode dispersion, has been proposed, but has not found widespread use today [37]. Such advanced monitors may become important as greater automation is introduced. In fact, there is an energy tradeoff between system automation versus manual diagnostics and maintenance. Network operators often have a "network operations center" (NOC) or other management office(s) with maintenance, planning, and operations staff. Accounting for the carbon footprint of these operational functions may help to further motivate the current trend toward greater automation. Increasing monitoring and automation in the line system, however, will need to be considered carefully from an energy perspective to achieve an overall efficiency benefit.

B. Line System Energy Trends

The various power consumption values mentioned above present numbers that might be typical for 10 Gb/s per wavelength, which is predominately used in core networks [7]. Today, line cards with up to 100-Gb/s line rates are available and, within the same transmission reach, can achieve higher energy efficiencies compared to 10 Gb/s. However, the power of the transmission (line) hardware, i.e., $P_{\rm CC} + P_{\rm LC}$, on the board is much larger. For 10 Gb/s, the transmission (line) hardware is roughly equal to the overhead or common hardware such as board controllers. Fig. 3 shows power consumption for line cards of different capacities that are representative of commercial equipment today up to 40 Gb/s. Noting the power law trend in the line hardware power, we extend it to 1 Tb/s following [38]. In practice, the exact quantities are dependent on many factors such as the client side configuration, and it is not clear that this trend will continue. For the common or conditioning equipment power we assume a constant component of 20 W and take the client modules and the board power losses to each be 10% of the line hardware. As the capacity increases, new optical or electrical hardware might be introduced such as dispersion compensation, optical preamplification, PMD compensation, balanced receivers, or stronger FEC. The line-component-power-dependent contribution ($P_{\rm LC}$) does not change substantially, but the conditioning components $P_{\rm CC}$ come to dominate the line card power. With this trend, the power of a single line card will exceed 1 kW for 1-Tb/s per wavelength. Furthermore, for 10 Gb/s and lower, the total line card power is weakly dependent on the line rate, but at higher rates, the line card power becomes dominated by the line hardware power and shows a correspondingly strong bitrate dependence. This behavior is not known to be fundamental and methods to reduce the slope will help facilitate scalable network growth in the future.

In addition to increasing the bitrate per wavelength in a line card, greater efficiency can be realized by using integration at the board and photonic levels. Considering a shelf that holds 15 line cards plus one 50-W controller card, integrating three 10 Gb/s line cards together to share the common equipment would reduce the shelf power from 1 kW to 750 W (including the 20% cooling loss). As the bitrate increases, however, if the transmission hardware follows the trend in Fig. 3, this advantage will

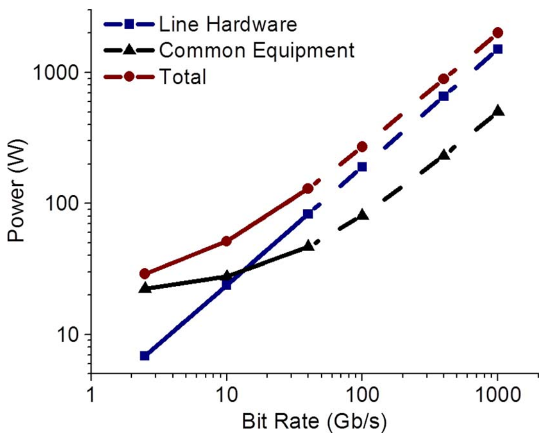

Fig. 3. ***Capacity-dependent optical line card power. Values beyond 40 Gb/s are extrapolated following the logarithmic trend.***

no longer be available since the common equipment will only be a small fraction of the total power. Efficiency might still be found at the board or chip level by sharing resources. For example, a single FEC chip might be used to process the signals from multiple wavelengths. In the limit of line hardware constrained power, efficiency improvements must come from the transmission system and component design, not capacity alone.

In a central office environment, the power is constrained by both the electrical feed limits and by thermal density limits. Assuming line cards today are at or near the thermal density limits, increasing the capacity of a line card following the trends in Fig. 3 will require a larger spatial footprint or card slot size in the shelf to avoid further increasing the thermal density. Given that central office space is limited, this will create an increasing thermal density bottleneck on capacity growth. Standards such as the Network Element Building System Telcordia GR-63-CORE limit the amount of heat that can be generated per equipment floor area. Telcordia standards, which are widely used in North America, suggest a thermal density per shelf for natural convective cooling of 740 W/m^2/m of vertical shelf space and 995 W/m^2/m for forced air cooling. Similar standards are followed in other parts of the world, such as the European Telecommunications Standards Institute (ETSI) for the European Union. Note that stacking two 1-kW shelves in a square meter area would exceed the Telcordia standard by a factor of two.

V. SPECTRAL EFFICIENCY AND ENERGY EFFICIENCY

Because energy has not historically been a primary design constraint in transmission systems, one might expect that there are opportunities for improving efficiency. Spectrally efficient modulation and multiplexing, which utilize the orthogonality of optical signals in time, frequency, polarization, quadrature, and space, play a central role in modern optical transport systems [39]. Increasing spectral efficiency has contributed to increasing energy efficiency due to a number of factors including the large common equipment and conditioning component power in commercial systems as described above and also the fact that the optical hardware could support more capacity. WDM channel widths are 50 GHz or wider, whereas until recently data rates were less than 40 Gb/s. Thus, by deploying transmission systems employing more efficient modulation formats, single channel data rates could be continuously increased without additional line system hardware or other modifications. Furthermore, an erbium-doped fiber amplifier spectrum can support up to approximately 90 WDM channels at 50-GHz spacing for long-haul transmission. The same basic amplifier has been used predominately since the first WDM systems. Once the EDFA band is full, then further traffic growth will drive the need for additional bandwidth and this will mean additional hardware: amplifiers, switches, multiplexers. The situation is analogous to passengers filling a bus. Until the bus is full, additional passengers can be accommodated with little or no increase in energy. Once the bus is full then additional buses or double-decker buses are needed to support the growing number of passengers—leading to an increase in energy use.

The past decade saw the rapid development of advanced modulation formats and multiplexing schemes in optical transport systems, particularly to support very-high-capacity transmission systems. However, improvements in spectral efficiency may come at the expense of increased hardware complexity and power constraints, contributing to the trend in Fig. 3. The complexity of the hardware, such as the transmitter/receiver structure and electronic processing, increases with the number of bits carried by each modulation symbol [27]. As such, supporting spectrally efficient modulation formats is likely to increase the power consumption of the transceiver hardware. The energy efficiency on a per bit basis needs to be evaluated for each modulation format and hardware implementation.

Returning to the Shannon capacity limit to examine the tradeoff between the spectral efficiency and energy consumption, the spectral efficiency (SE) is defined as the number of bits per symbol [40]

$$\mathrm{SE} = \frac{\log_2 M}{N/2} \tag{7}$$

where M and N are the modulation level (number of constellation points) and the signal dimension (e.g., quadrature, polarization), respectively. For an optical signal transmitted with an average power of P and a bitrate of Br,

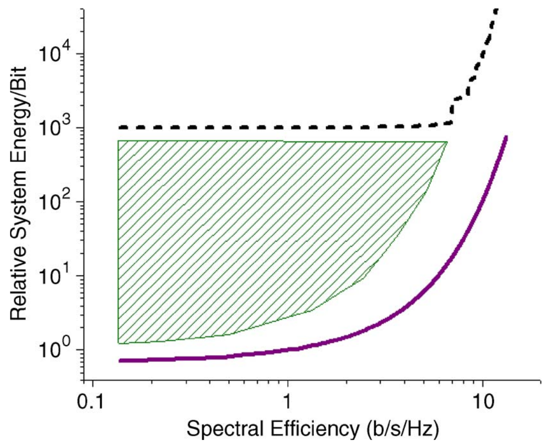

Fig. 4. ***System efficiency (dashed) and the ideal Shannon limited efficiency (solid). Hatched area shows region for potential efficiency improvements.***

the average energy per bit is given by $E_b = P/Br$. In a coherent optical system the additive white Gaussian noise is assumed to have a variance of $N_0/2$. The Shannon capacity limit relates the spectral efficiency and the average transmitted energy per bit for such a channel as

$$\frac{E_b}{N_0} = \frac{2^{\mathrm{SE}} - 1}{\mathrm{SE}}. \tag{8}$$

Thus, increasing the spectral efficiency requires an increase in the energy per bit. Fig. 4 shows a plot of the system energy per bit in units of HN_0 for a Shannon limited AWGN system (solid), where H is a system-dependent gain factor relating the received energy to the transmitted energy [14]. Also shown is the total system energy from (1) and (2) (dashed). The energy per bit is multiplied by the system inefficiencies to get the LC terms and then the other elements that are independent of E_b are added, creating a system power floor apparent at low spectral efficiencies [14].

Furthermore, recent studies [41] have shown a relationship between the maximum transmission reach and the spectral efficiency; systems employing high spectral efficiency signal formats require a higher SNR, which in turn leads to shorter spans between optical amplifiers or a shorter overall distance between regeneration points. Thus, very long reach transmission systems favor the use of more robust but less spectrally efficient modulation formats, and require higher speed modulators and electronic circuitry. Conversely, shorter reach systems can take advantage of more spectrally efficient modulation formats, achieving higher capacities within a given wavelength channel.

Recent research that addresses transmission system energy issues falls mostly into two categories: 1) evaluating the impact of modulation formats on the optical transport systems [27]; and 2) designing modulation formats that offer a good tradeoff between spectral and power efficiency [40], [42], [43].

In [27], Tucker analyzes the influence of the modulation formats on the combined per bit energy consumption of three key components of a point-to-point transport system: the transmitters, the receivers, and the optical amplifiers. In particular, the minimum energy per bit quadrature amplitude modulation (QAM) channel is compared with that of on–off keying (OOK) and binary phase shift keying (BPSK) in three regions: the region in which transmitters and receivers dominate the total energy, the region in which amplifiers dominate, and the region in which both transmitters/receivers and amplifiers contribute approximately equally, respectively. The comparison leads to the following key results.

1) When transmitters and receivers dominate the energy consumption of the transport system, the total energy per bit of a QAM system is likely to be larger than that for OOK or BPSK. This is due to the increased energy consumption incurred by the modulator and driver of a QAM transmitter and analog-to-digital converter (ADC) of a QAM receiver.
2) When amplifiers dominate, the minimum energy per bit increases by 2 dB for each doubling of constellation size.
3) When transceivers and amplifiers contribute equally (as a result of optimizing the repeater spacing), the per bit energy consumption of a QAM system is $\sqrt{2\mathrm{SNR}_{\mathrm{bit}}}$ higher than that of an optimized BPSK system.

Some recent research [40], [42], [43] focuses on finding modulation formats that achieve the minimal transmission energy per bit. In particular, for a given signal dimension N, and a constellation size M, and asymptotic SNR, numerical optimization techniques are used to find the best constellation that has the highest asymptotic energy efficiency. In [40], by computing the best sensitivity for 4-D modulation formats up to 32 levels, Karlsson and Agrell identify polarization switched QPSK (PS-QPSK) as the most power-efficient modulation format for uncoded coherent optical systems, which is shown to have an asymptotic gain relative to BPSK of 1.76 dB.

The total system energy in Fig. 4 assumes a 2000-km reach system, includes the amplifier line card power, and multiplies the transceiver line card power by a factor to account for the regeneration required by the spectral-efficiency-dependent maximum reach [41]. Thus, the total system energy curve increases in steps at high spectral efficiency due to the multiplicative regeneration requirements. The Shannon curve provides the lower bound for error-free transmission. Depending on the specific system details such as distance and capacity, there exists a region, such as that indicated by the hatching, within which efficiency improvements will be possible, and this region should be a focus for research and development by transmission system engineers. There is an offset between the hatched region and the Shannon limit due to technological and practical constraints on system reach and minimum power. Developments to improve energy efficiency together with spectral efficiency would involve a wide range of technologies, modulation formats, and channel coding, optimized across the system.

The study of energy-efficient modulation formats necessitates a comprehensive evaluation (in terms of energy per bit) of the combined effect of signal generation/recovery in the electronic domain, signal modulation/demodulation between the electronic and optical domain, and signal propagation in the optical domain. Issues such as electronic processing, receiver structure, optical noise, dispersion management, and nonlinear effects need to be addressed jointly and are still largely open topics for research.

VI. OPTICAL NETWORKS

In seeking to improve the energy efficiency of data networks, several opportunities arise. These span a range of areas from network architecture design to the management of network equipment. Some of these will require development of new features in switching and transmission equipment. First, we consider methods to improve the optical network efficiency for a given architecture, which primarily involves improving the equipment utilization. Next, we consider energy use from a network design perspective and describe the many optimization dimensions available. Finally, we discuss optical switching and transmission system issues that must be addressed when considering these network level energy optimizations. The different architectural options and design tradeoffs described here are listed in Table 2 along with the corresponding document section and bibliographic references.

A. Energy-Efficient and Green Network Techniques

In general, network equipment is provisioned with capacity that exceeds the mean traffic. This overprovisioning is used to accommodate traffic bursts, provide spare capacity for use in the event of a failure (protection), and allow for mean traffic growth over a period of time. Near the edge of the network, with less aggregation, traffic bursts are larger and more common, therefore the overprovisioning ratio can be quite large, sometimes as high as a factor of 100 times. In the core network, the traffic is more uniform and the overprovisioning ratio can be as low as a factor of 2–4 times. Introducing dynamic functionality in the network thus has the potential to reduce the total equipment energy consumption by a factor of 2–100. During periods of low traffic levels, some equipment could be turned off to save energy. As the traffic demands vary across the network, resources would be redirected along new network paths to follow these changes.

Table 2 Energy-Efficiency Architectural Options and Design Tradeoffs for Optical Networks

Category	Option or Design Trade-off	Sec.	Ref.
Network Arch.	Shutting off and overprovisioning of networks vs. rate adaptation or load-proportional operation	VI. A	[44][45]
	Renewable energy use vs. energy efficiency of optical transport equipment	VI. A	[31][50]
	Dense vs. sparse direct optical connections	VI.B	[51]
	Transport vs. storage energy efficiency	VI. B	[30][53 – 56]
	Transport vs. processing energy efficiency	VI. B	[57]
	Blocking performance vs. optical switch size	VI. D	[75][77]
Optical Trans.	Transmission margin vs. performance	VI. C	[58 -60]
	Data rate vs. transmission reach	VI. C	[4] [61 - 68]
Optical Switch.	Dynamic wavelength switching vs. control overhead	VI. A & E	[30][46] [47]

The ability to benefit from dynamic network capabilities is limited by many factors. Scaling the energy use of devices continuously with traffic load may not be possible in some cases. The energy efficiency of many devices is derived only from its ability to transfer between "on," "idle," and "off" states. Thus, the savings will only come from the extent that devices or network elements can be placed into a low energy "idle" state or a zero energy "off" state. The speed at which devices can change state and adapt to traffic variations is an important factor. If the speed is too slow to follow the variations, then overprovisioning will still be required. Short-term traffic bursts are often unpredictable and rapid. Depending on the equipment or components being shut down, there may also be energy costs in re-enabling the equipment and reconfiguring the network to take account of the configuration changes.

Overly aggressive shutting off and provisioning of network equipment can degrade the network performance [44]. On the other hand, long-term traffic growth, diurnal, and seasonal changes occur on long time scales and often follow regular patterns that might allow for the implementation of smooth and predictable transitions between network equipment energy states. In these situations, rate adaptation or load-proportional operation might be preferable to sleep mode techniques. When energy consumption of a device can be scaled with traffic load there is less benefit to transitioning to an "idle" or "off" state and the break-even point depends on the network topology [45].

Introducing dynamic capabilities also presents the possibility for unstable behavior or long settling times [46]. Instabilities have been studied in the physical-layer optical power [47] and in higher layer control planes. A physical layer that responds to traffic changes would require cross-layer control that may further complicate these already complex dynamical systems.

Other issues can also impact the efficiency of dynamic operation. For example, when switching wavelengths in response to flows or file transfers, a large variance in the flow or file sizes can lead to congestion. This can be modeled as an $M/G/1$ queue with the first-in–first-out (FIFO) model. The average file transfer time T depends on the capacity C, the average file size $E[B]$, the variance of the file size $\mathrm{Var}[B]$, and the download rate R as follows [48], [49]:

$$T = \frac{E[B]/C}{1 - E[B]R/C}\left(1 + \frac{E[B]R}{2C}\left(\frac{\mathrm{Var}[B]}{E[B]^2} - 1\right)\right). \tag{9}$$

Thus, as the variance of the file size increases, the download time increases and the energy efficiency degrades. The wavelength capacity can only be used to transfer one file at a time. As a result, small file transfers are likely to be suspended by an existing large file transfer. The increase of the queuing delay can translate into extra power consumption of the queuing hardware. As such, it is suggested [48] that dynamic wavelength switching better suits applications in which the transferred file sizes do not vary significantly, such as video-on-demand (VOD) services.

Rather than focusing only on efficiency, the carbon footprint of networks can be improved by focusing on the use of renewable energy sources. Renewable sources tend to be volatile, but an approach that has gained interest in the data center community is referred to as "follow the sun, follow the wind." This strategy involves moving content between data centers or between network elements in order to access renewable energy when and where it is available. A related strategy involves placing centers at locations where natural cooling is available, reducing the need for cooling and ventilation which consume an appreciable proportion of the energy used in data centers and major network centers [31], [50]. Importantly, these strategies take advantage of the high energy efficiency of optical transmission systems in order to enable the more power hungry network elements to maximize the use of renewable energy.

B. Network Architecture Transmission, Switching, and Hosting

Several different energy tradeoffs arise in designing an energy-efficient network architecture. In general terms, an efficient architecture strikes a balance between the scale of network transmission resources, switching resources, and the concentration of content and computing resources.

A clear example of such a tradeoff is seen in the design of an optical network connecting major network nodes and data centers. Providing a greater number of direct connections between these nodes reduces the level of traffic handled by the IP routers and the number of router ports required, but requires the addition of more transmission capacity. Since IP routers can be 1–2 orders of magnitude less efficient than the transmission hardware in a link, on an energy per bit basis, the network efficiency can often be improved by increasing the number of direct optical or transparent connections between nodes. In addition, router ports are generally more expensive than optical amplifiers and transponders, thus there are potential savings both in cost and energy. However, taken to excess, the provision of large numbers of direct optical connections may require substantially more investment in cable and transmission infrastructure, negating the cost savings [51].

Network availability also needs to be considered in this optimization, because with more direct connections, a cable cut can cause the loss of a greater number of IP links. This in turn means that more restoration resources need to be provided.

In addition, transport utilization plays a role here. Higher utilization of the optical transmission lines will lead to more efficient optical systems, but may cause significant congestion in the higher layers resulting in a less efficient network overall [52]. Designing for overly high utilization can also mean that less reserve capacity is available for service restoration.

In the case of content delivery networks, placing content closer to the end user can reduce the number of transport hops used to deliver the content. In core networks, the relationship between the transport hop count to a replication site and the number of corresponding replication sites can be shown to roughly follow a power law relationship of the form $H(n) = A(N/n)^{\alpha}$, where N and n are the number of nodes and replications, and A and α are constants that characterize the network topology [30]. Increased content replication, however, leads to greater storage energy and cost. This tradeoff has been studied for both content distribution networks [30], [53], [54] and content centric networks [55], [56].

Source coding, used to compress a data stream, reduces the number of bits transported in the network, but this comes at the expense of greater processing power consumption. For the case of software-based compression using servers or PCs, uncompressed data transmission ($\sim 10^{-7}$ J/b for ten core hops) was shown to be more efficient than compressed data transmission after including the compression energy ($> 10^{-6}$ J/b depending on the compression ratio [57]). However, for content that is accessed by many users over a period of time, an investment in compression can be worthwhile.

C. Transmission Dependent Network Optimization

Other network energy tradeoffs are sensitive to the optical transmission design. From a network perspective, transmission systems create a network of lightpaths which can be viewed as a bidirectional, directed, and weighted graph. The weights contain different attributes such as capacity, quality of service (QoS), reliability, cost, or energy. The routing and assignment of wavelengths along these lightpaths is referred to as the RWA problem and can include constraints such as transmission impairments or energy consumption [58]–[60]. A key tradeoff in network design is the balance of transmission margins against transmission performance optimization or control through either offline or online techniques. Performance optimization or control includes RWA using more accurate constraints or optical performance monitoring and impairment compensation. Larger margins use more equipment (e.g., more frequent regeneration), whereas smaller margins increase the likelihood of faults or failed operations, which then increases the maintenance and operational energy consumption. Increased monitoring and control may likewise increase energy consumption.

In [41], Winzer compares the energy efficiency of a parallel transport system (e.g., spatial multiplexing, multicore or multifiber transmission) with that of a single-mode system using multilevel modulation to achieve equivalent capacity. Since each of the parallel signals is transmitted at a lower data rate, longer transmission distances are achieved without electronic regeneration. In this way, a parallel system has the potential to achieve energy savings in comparison with a single-channel system—potentially two orders of magnitude improvement in energy efficiency for a system supporting 20 b/s/Hz over 1500 km has been shown, depending on the relative energy cost of the additional transceivers in the parallel case.

Much recent research focuses on mixed-line rate [61]–[63] and bitrate adaptive [64]–[67] approaches in addressing the tradeoff between data rate and transmission reach. In the mixed line rate approach, transponders with different data rates are deployed. Low data rate transponders are used for demands that span long transmission reach, while mid or high data rate transponders are used for short reach demands. The bitrate (flexible) adaptive approach provides more flexibility in data rate. By using optical orthogonal frequency division multiplexing (O-OFDM) for high spectral efficiency in analogy with RF systems, the data rate can be adjusted in response to the change of bit error rate, impairment, and traffic demand. The primary energy benefit of mixed-line rate and bitrate adaptive approaches comes from the reduction in the numbers of optoelectronic regenerators needed on transmission links. For example, in a 28-node core network with a total traffic load of 15 Tb/s, a bitrate adaptive approach was shown to save about 40% in the number of transceiver line cards required, in comparison to using a fixed-line rate of 40 Gb/s [68]. These adaptive capabilities are expected to incur an additional cost in terms of energy, complexity, and expense, thus the benefits must always be evaluated against such potential offsets.

D. Optical Switching

Switching in networks can take place at different physical layers using a variety of technologies, and over different time scales. Circuit switching is used to set up a path between network nodes that is continuously available, generally for a period which might range from seconds to months or years. Circuit switching in the optical layer involves preallocating resources to provide one or more wavelengths for the duration of the connection. Electronic circuit switching may route traffic at the subwavelength level using reallocated time slots in a time-division-multiplexed (TDM) stream, for the duration of the connection. In contrast, packet switching or flow switching directs individual bursts of traffic along particular paths; these traffic bursts may last from nanoseconds to milliseconds. In this case, resources are allocated at the level of individual packets or bursts.

Fig. 5 shows the range of equipment that might be used to achieve these different forms of traffic switching and processing functionality, together with typical energy per bit efficiencies for 2009 products, and projected values for 2015.

1) Optical Path Switching: Wavelength-based switching in commercial systems today is used to provide turn-up "circuit switching" in which a dedicated communications channel is set up before the end users communicate and then the connection is fixed over a period of months or years. Subwavelength (usually electronic) switching can provide both circuit and packet switching. With packet switching, data flows are segmented into packets which are separately multiplexed and switched to share the available communications channels.

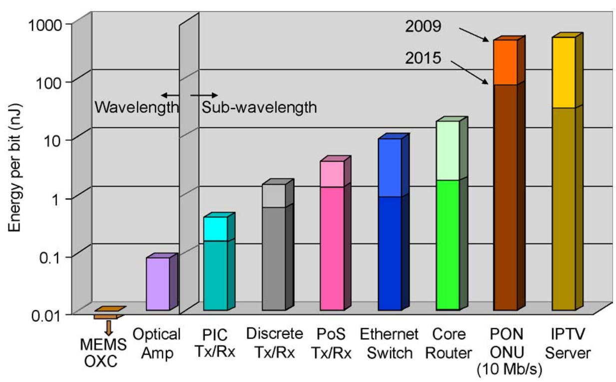

Fig. 5. ***Energy per bit of different types of network elements. The elements are divided into those that operate at the electronic (subwavelength) level and those that operate on wavelengths. The lightened tops of the columns represent the improvements in energy efficiency expected for these elements over the period 2009–2015. PoS = packet over SONET. PIC = photonic integrated circuit. Tx/Rx = transmitter/receiver. OXC = optical cross connect. PON ONU = passive optical network optical network unit. (Source: [7].)***

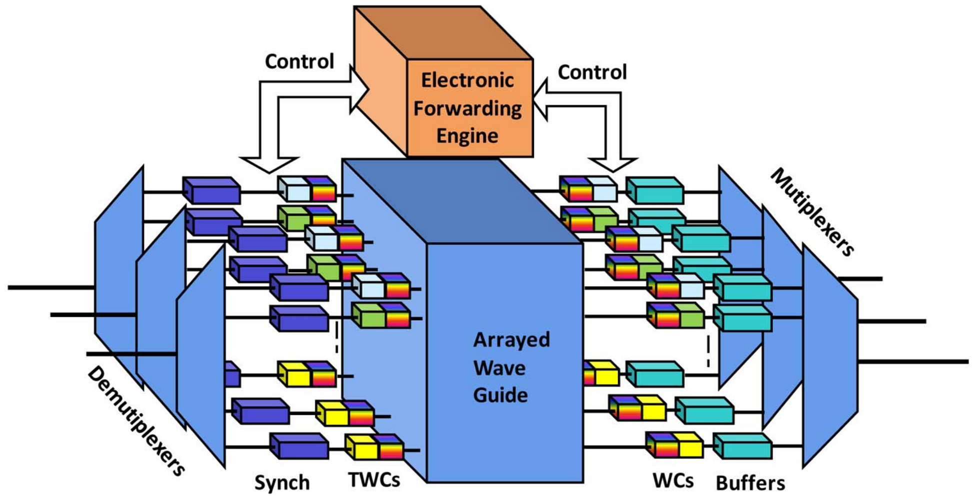

Fig. 6. *Optical router architecture consists of an electronic forwarding engine to electronically process the packet headers and control the tunable wavelength converters (TWCs) so that the incoming packets are steered to the appropriate output by the AWG. The packets output from the AWG are then wavelength converted (WC) to the correct wavelength for multiplexing into the output fibers. The router may also require synchronization at the input and buffers at the output to avoid packet collisions. The packet payloads remain in the optical domain. (Source: [74].)*

Wavelength switching elements are the most energy efficient because they do not process traffic on a packet-by-packet basis. The energy per bit increases for sub-wavelength switching with decreasing aggregation and increasing network-layer functions.

Optical switching can be implemented using technologies such as microelectromechanical systems optical cross connects (MEMS OXC), which can provide switching on a time scale of milliseconds at very low energy per bit (roughly 0.1 nJ/b) [69]. These devices redirect optical channels without any form of processing the data within those channels and thus can only be used as a switch fabric or for circuit switching functionality. Since the switching energy is independent of the channel capacity, the efficiency is directly proportional to the data rate in the channel being switched.

Often the power of an optical switching device is dominated by the temperature control or other conditioning elements needed to maintain the switch state over long periods of time. Computations required to determine the switch configuration and the associated control protocols also need to be taken into account and may dominate the power consumption.

2) Optical Packet Switching: To implement optical packet switching and routing requires the adoption of high-speed optical switching technologies. Many different methods of incorporating optical capabilities into packet routers have been studied in the literature [69]. A widely adopted efficient approach is to modulate the packet data onto the output of rapidly tunable laser sources and rely on a passive arrayed waveguide grating (AWG) to steer the signal to the appropriate output port [70]–[73]. The optical router consists of an electronic control/forwarding engine, which examines incoming packet headers, using that information to redirect the packet payloads to the appropriate output port; the payloads themselves remain in optical form. (See Fig. 6.)

Recent studies found that these AWG (and other) types of optical routers do not provide significant energy savings relative to an electronic router [69], [74]. This comparison is without including the power consumption of monitoring required for fault detection, location, and management of the (unregenerated) optical payloads that will be required in the optical router, further reducing the potential efficiency of the optical solutions. In fact, only the switching and buffering functions were considered.

The power consumption in routers used today is dominated by forwarding engines, which implement packet inspection functions such as pattern matching for address resolution, QoS, and firewall applications. These are complex high-speed computation-intensive operations for which energy-efficient implementations using optics are not known. Laboratory demonstrations of optical packet switching do not include these advanced functions. Thus, in most cases, a laboratory optical packet switch cannot simply be compared with a commercial router, which is optimized to provide these advanced functions.

3) Optical Switch Architecture and Technology: The energy efficiency of a switch will depend on its architecture and operation. Efficient binary switching or Benes fabrics can achieve switching energy that scales with the logarithm of the number of ports. For example, a $1 \times N$ binary switch is composed of a tree of 1×2 switches of which $\mathrm{Log}_2 N$ switches are used to create a unique connection between the input port and any of the N output ports. For clocked switching gates, common in electronic fabrics, switching energy is dissipated with each bit level change on each clock cycle. Furthermore, additional energy is consumed distributing the clock signals. Clockless logic can be implemented to avoid much of this additional energy dissipation.

Optical circuit switches are often implemented as clockless analog switches that dissipate minimal excess switching energy for each bit transition. If the overhead energy (for example, in the control electronics) is sufficiently small, then the switching energy can be made to depend on the frequency at which switching transitions occur. For circuit switching, such changes might be measured in seconds, hours, or even days. In contrast with flow or packet switching, switch events may occur on time scales of nanoseconds to milliseconds. For switching at the nanosecond or shorter time scales, technologies such as MEMS are not appropriate. Instead, optical switching on these time scales is typically based upon fast optical nonlinear phenomena such as the optical Kerr effect, and three and four wave mixing. Devices based on these phenomena almost always require a continuous supply of optical or electronic power [75], [76], which significantly increases the energy per bit.

A further potential limiting factor for the energy efficiency of optical switching is the need to resolve wavelength conflict (blocking) during network switching operations. Without using wavelength conversion, one can always reduce the probability of wavelength conflict by increasing the number of wavelengths and ports [77]. However, such a "no wavelength converter" approach does not scale well with the network size and traffic demand. As optoelectronic regeneration remains the most practical approach for wavelength conversion, the energy benefit of optical switching diminishes as the percentage of ports that are equipped with regenerators for wavelength conversion increases. In [78], it is estimated that, with a port count of 1000, every 25% increase in the number of wavelength converters increases the power consumption of an optical switch approximately twofold in comparison to the consumption with no wavelength conversion. In this regard, further tradeoffs need to be evaluated between energy consumption and switched network performance [44].

E. Transmission Issues in Dynamic Networks

Using transparent optical switching to realize the network-wide efficiencies described above requires not only the switch functionality but also support for wavelength switching within the optical transmission system, which may involve propagation over thousands of kilometers. As with many transmission technologies, the difficulties related to transparent optical switching increase with transmission distance and bandwidth (or number of channels). Rapid, stable switching has already been demonstrated in small networks under various constraints [73]. Achieving this functionality in long-haul or even large metropolitan networks is still an open research problem. Introducing dynamic functionality in the physical layer also complicates the operation of higher layers; e.g., disabling a link within a network may cause routing information updates to be advertised around the network and routing tables updated. Rapid fluctuations in network utilization can cause transport protocols such as TCP to exhibit nonlinear and unstable dynamics. If both the physical layer and the higher layers are making independent decisions about resource utilization, then the potential can arise for competition and instability. Cross-layer protocols can enable coordination, but may further complicate the design of stable algorithms. Control plane capabilities have been demonstrated recently with the potential to achieve rapid reconfiguration in the higher layers, but the physical-layer aspects have not been addressed [79].

Dynamically reconfigured networks place some additional constraints on the design of the underlying transmission system layer. In general, there are many different controls in transmission systems that need to be tuned in order to achieve error-free transmission. These controls take time to find a new steady state after a network change [80]. When a wavelength is removed or added, different controls will respond on different time scales. Fast constant gain control in the amplifiers will adapt on time scales of microseconds to milliseconds. At the network level, variable optical attenuators, amplifier mean gain and tilt, wavelength-selective switch channel attenuators, and dynamic gain equalization filters are all adjusted on much slower time scales. When individual channel power levels are involved, then channel monitors are needed to determine the power levels. A channel monitor may take 1 s or longer to complete a scan and multiple iterations may be required. Such tuning adjustments can impact the downstream power levels. In constant gain controlled amplifiers, which are most commonly used, the channel power fluctuations are coupled through the wavelength-dependent amplifier gain and nonlinear effects in the fiber [81]. Therefore, downstream from a channel-power tuning adjustment, both the channel in question as well as other channels on the path can all be impacted. These interactions have been shown to lead to instabilities when adjustments are taken in parallel [47]. Serial adjustments or slow, small steps can dampen these effects, but also significantly stretch the time required to retune a network after a wavelength switching event. These network control and other provisioning delays can impact the overall efficiency of the switching functionality.

The efficiency impact of network control delays can be quantified by considering the ratio of the total quantity of data to be transferred B in the new configuration to the control delay-bandwidth product T_0C, where C is the bandwidth of the wavelength channel and T_0 is the full setup and turndown control delay time. The efficiency in energy per bit can be written [30]

$$e_{\Delta\lambda} = \gamma e_{\lambda} \tag{10}$$

where $\gamma = (B + T_0 C)/B$, and e_λ is the efficiency (energy per bit) of the established wavelength connection. Thus, wavelength switching is only efficient if the volume of data is much larger than the control delay-bandwidth product. As shown in [30], this break-even traffic volume size is proportional to the control delay-bandwidth product scaled by the ratio of the WDM equipment to router energy efficiency. With current technology, T_0 is on the order of minutes and the efficiency of WDM equipment can be an order of magnitude better than that of routers in a long-haul system. If the wavelength capacity is 10 Gb/s, dynamic wavelength switching is only beneficial for delivering a file on the order of 100 Gb. Note that the control delay might be reduced through introducing new controls or reducing the transparent reach of the network. The added hardware, however, increases the power consumption of the network. Indeed the control delay is strongly dependent on the size of network (in terms of either distance or number of amplifiers/elements) and the capacity (both in terms of the total system capacity supported and the number of wavelengths involved in the switching event).

Establishing a new wavelength path also requires careful network planning in order to determine whether the new path will satisfy the transmission performance and margin requirements of the system. Planning tools are used offline to determine if a new configuration will be successful. For wavelength switching, real-time tools would be needed to make these performance estimations. If the estimations are aggressive and result in numerous errors, then the new configurations will fail and result in additional switching to find alternate connections. This further compounds the control delay and thus reduces the network energy efficiency. On the other hand, a conservative estimation may limit the range of configurations that can be used for switching and thus also limit the efficiency.

VII. ENERGY OF A SERVICE

The Internet provides a diverse range of services and network use is determined by these services, including, for example, telephony, video, e-mail, web browsing, and search. As the issue of global warming grows in significance, end users are becoming interested to know the carbon footprint of particular services (e.g., a video conference service) rather than the carbon footprint of a generic broadband network connection. Services impose different requirements on the network, which in turn dictate the

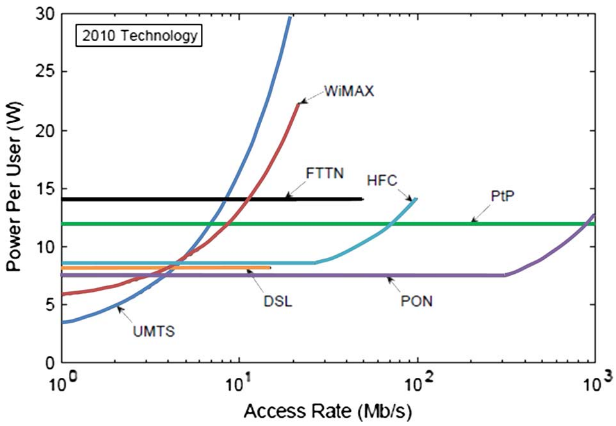

Fig. 7. ***Power per user for access network technologies as a function of access rate. Although wireless access is relatively energy efficient at very low rates, it becomes the least efficient as rates increase. PON provides the most energy-efficient access technology for rates less than around 1 Gb/s. Note that these comparisons are based upon 2010 technologies.***

network energy use. Networks have benefited from transport infrastructure, including routers, switches, and transmission hardware, which is largely service agnostic. However, increasingly unique functionality is being introduced into transport hardware in order to support the disparate service requirements and this comes at a cost in terms of energy consumption. Future networks will need to reconcile these service-specific requirements against their energy consumption.

Today, there is a variety of ways by which we can access the Internet and other ICT networks to utilize a service. These include: PONs, wireless access, fiber to the node [with digital subscriber line (DSL) from the node to the home], hybrid fiber coaxial cable, and point-to-point Ethernet. Fig. 7 shows the power consumption per user for a range of access network technologies [4]. We note that as access speeds increase, wireless access [via the Universal Mobile Telecommunication System (UMTS) or Worldwide interoperability for Microwave Access (WiMAX)] becomes the most power consuming while PON is the most power efficient. These results represent a challenge to the telecommunications industry because mobile access is the most rapidly growing access technology due to its ability to provide "anywhere" access to services [5], [82]. Many point to the convergence of optical and wireless networks as a potential solution, where convergence refers to techniques such as RF over optical [83] (sending analog RF signals from an antenna or base station directly over an optical fiber to a central processing unit), or optical base-station networks and the use of fiber all the way to the base station or even the antenna. These techniques may lead to complex optical networks as mobile networks move toward small cell architectures with cooperative and autonomic functions.

Services differ in the requirements they place on the network in terms of traffic volume, speed of delivery, security, QoS, latency, and jitter. This in turn leads to the use of different content deployment architectures, different transport protocols, and the need for routers to be aware of and give priority to certain types of service traffic. When a service is concentrated in just a few data centers worldwide, there can be energy savings in the storage and computing needs at the data center, but with energy cost in the transport of requests, data, and control packets to and from the user. Conversely, as described in Section VI, there are services which are dispersed among a large number of data centers, delivering content with low latency and low transport energy cost, but the very replication of data among so many sites could raise the total energy cost of the service. As a general principle, content which is more popular, frequently accessed, or primarily of regional interest should be replicated close to the user, while infrequently accessed materials need to be stored in just a few locations, reducing the storage cost but incurring a higher transport cost [29], [56]. Recently, ideas such as content-centric networks and nano data centers have been proposed to improve the energy efficiency of content distribution through optimal content placement [55], [56].

The energy consumption of services provided via the Internet has been studied for several cloud-based services [53]. Cloud-based services provide end users with access to high powered processing and storage in facilities located within the Internet on an as-needed basis, avoiding the need for individual users to invest in, maintain, and power their own servers. Looking at the cloud from an energy viewpoint, it is possible to spread the power consumption of these facilities across many thousands of users and so dramatically reduce the power per user to improve overall service energy efficiency. Recent results show that cloud services are not always as energy efficient as undertaking tasks on a desktop PC. This is particularly so if there is a substantial amount of transport of data between the end user and the cloud facility (which may be located in another city or even another continent), or transport of data between cloud facilities [53]. Therefore, reducing the transport power consumption by increasing network transparency, i.e., using optical bypass techniques to reduce electronic processing, can help to improve the efficiency of transport-intensive cloud applications.

VIII. LEGACY SYSTEMS AND TRANSMISSION SYSTEM ARCHITECTURES

Bringing together the threads of the discussion so far, we see that the current "state of play" of energy consumption in telecommunications networks is a complex interplay of technologies, equipment design, architectures, services, customer preferences, and much more.

The history of telecommunications also plays a very significant role in our current predicament. Before the growth of the Internet, telecommunications networks were dominated by voice traffic and so were designed to accommodate traffic that must arrive in real time, in serial order with minimal errors. In contrast, during the early years of the Internet, packet traffic did not need to satisfy any of these requirements. However, this is now changing.

This history has resulted in network designs in which IP traffic was initially an "add-on" in that IP traffic was merely force fitted into the TDM protocols, such as SDH and SONET, used to transport voice traffic. This resulted in a "multilayered" network often referred to as IP/SDH (SONET) or, in some cases, IP/ATM/SDH (SONET). In addition, it was commonplace for path redundancy to be provided at each of these layers.

As IP traffic grew, it became obvious that these multilayered networks were becoming increasingly inefficient and expensive, and new, simple, packet-oriented network architectures evolved. But at the same time, as the range of Internet services broadened (to include real-time data and provide a measure of QoS,), more protocols have been added to the IP protocol stack to cope with these demands. Each layer of protocol requires more data processing and hence more power consumption.

Viewing Fig. 8, we see that those network elements that undertake heavy data processing are least energy efficient. With legacy TDM transport of IP packets, such as IP/SDH [Fig. 8(a)], all traffic entering the node is processed in the IP router, which has an energy consumption of approximately 10 nJ/b, regardless of whether the traffic is destined for that node.

By appropriately grooming traffic, we can enable IP traffic not destined for a node to bypass the IP router in that node [Fig. 8(b)], via a simple TDM circuit switched path, which may be next-generation SDH/SONET or optical transport network (OTN) and consumes the order 1 nJ/b. This also allows a reduction in the size (and port count) of the IP router to improve energy efficiency.

More recently, SDH/SONET has been modified to improve its ability to efficiently carry packet data (so called next-generation SDH/SONET). In addition, the OTN standard has also been developed to provide efficient transport of packet and TDM-based services [84].

In addition, with the introduction of the generalized multiprotocol label switching/automatically switching optical networks (GMPLS/ASON) automated control plane, the use of bypass at multiple levels in a node can be automated to optimize the network for today's Internet with respect to energy consumption, as depicted in Fig. 8(c) [85].

Network design for the future needs to make the appropriate use of all of the technologies available in order to evolve toward supporting future services, and include energy along with the traditional metrics of cost and

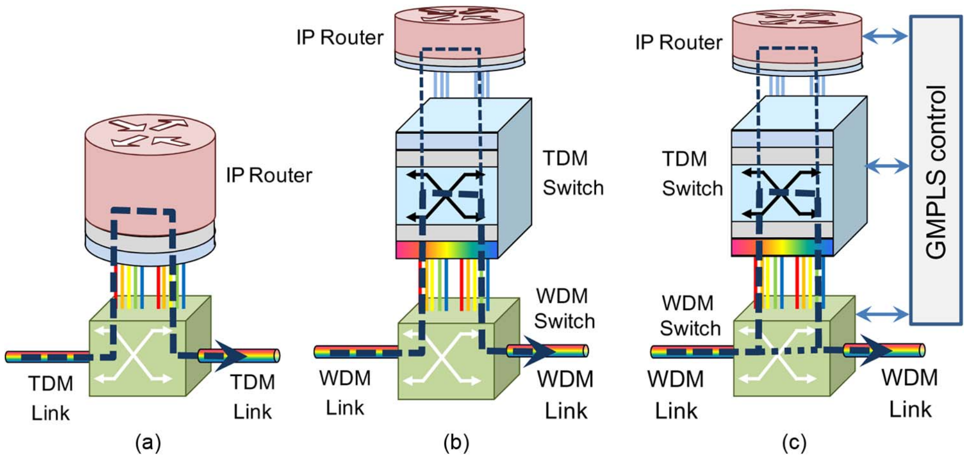

Fig. 8. *Examples of network layers. (a) IP over TDM transport such as IP/SDH (IP over SDH) required all packets to be processed within the IP router. (b) IP/TDM/WDM (IP over TDM over WDM) allows for some traffic to be processed within the TDM cross connect. (c) GMPLS/ASON automates the network allowing dynamic reconfiguration of the network in all the IP, TDM, and WDM layers. By shifting from packet switching in the IP layer to circuit switching in the TDM and WDM layers we can significantly improve network energy efficiency.*

operational complexity. GMPLS/ASON functionality is available in the latest commercial equipment, although it may take time to roll the legacy equipment over to the state of the art. Furthermore, this functionality today is implemented without wavelength switching. Optical networks potentially can use this capability to support greater efficiency through a dynamic physical layer provided that the transmission and control dynamics problems discussed above can be addressed.

IX. CONCLUSION

Energy is a focus today in the design of terrestrial optical transmission networks due to central office thermal constraints, and will become increasingly important in the future to ensure continued network scaling with demand, controlling network carbon footprint, and enabling eco-sustainable network applications. While energy has always been essential to the physics of communication, only recently have energy efficiency and carbon footprint become focal points of technology research and system design. Data networks in particular have matured over the past three decades to a point in which they are starting to encounter challenging physical constraints associated with energy and capacity. The key energy-related issues expected to shape future optical transmission networks are listed below:

- shift to parallel systems that require integration at the photonic, board, and system level;
- system and network optimization balancing energy use with other parameters such as reach, coding strength, and spectral efficiency;
- realizing energy-efficient dynamic physical-layer network capabilities;
- cross-layer network design for energy-efficiency and service-aware optimizations.

In this work, we have highlighted current work in these and related subjects and pointed to potential areas for continued progress. The energy consumption in typical commercial transmission systems was analyzed to provide guidance for future developments. Research on energy-efficient communication techniques is in its infancy and promises opportunities for innovation across all levels of technology from components to systems to networks. Optical networks are a foundational technology for modern communications. The challenge for the future is to devise new approaches to optical networking that provide energy-efficiency improvements at a rate commensurate with the growth in service demand. ■

REFERENCES

[1] N. Czernich, O. Falck, T. Kretschmer, and L. Woessmann, "Broadband infrastructure and economic growth," *Econ. J.*, vol. 121, no. 552, pp. 505–532, May 2011.

[2] European Commission, "The implications of ICT for energy consumption," Impact Study Final Rep. 09/2008, 2008.

[3] The Climate Group, *SMART 2020: Enabling the Low Carbon Economy in the Information Age*, Global e-Sustainability Initiative (GeSI), 2008.

[4] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, "Energy consumption in wired and wireless access networks," *IEEE Commun. Mag.*, vol. 49, no. 6, pp. 70–77, Jun. 2011.

[5] D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume, "Power trends in communication networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 275–284, Mar./Apr. 2011.

[6] T. Asami and S. Namiki, "Energy consumption targets for network systems," in *Proc. Eur. Conf. Opt. Commun.*, 2008, DOI: 10.1109/ECOC.2008.4729236.

[7] J. Baliga, R. Ayre, K. Hinton, W. Sorin, and R. Tucker, "Energy consumption in optical IP networks

J. Lightw. Technol., vol. 27, no. 13, pp. 2391–2403, Jul. 1, 2009.

[8] O. Tamm, C. Hermsmeyer, and A. M. Rush, "Eco-sustainable system and network architectures for future transport networks," *Bell Labs Tech. J.*, vol. 14, pp. 311–328, 2010.

[9] World Wildlife Federation, *The Potential Global CO_2 Reductions From ICT Use, Identifying and Assessing the Opportunities to Reduce the First Billion Tonnes of CO_2*, 2008.

[10] G. Cook and J. V. Horn, *How Dirty Is Your Data?* Greenpeace International, 2011.

[11] NTT Communications, *Green ICT*, 2011. [Online]. Available: http://www.ntt.com/business_e/feature/green_ict.html

[12] *Green House Gas Protocol*. [Online]. Available: www.ghgprotocol.org

[13] P. W. Huber, "Dig more coal—The PCs are coming," *Forbes Mag.*, 1999. [Online]. Available: http://www.forbes.com/forbes/1999/0531/6311070a.html

[14] D. Kilper, K. Guan, J. Llorca, G. Atkinson, and R. Tucker, "Coding and capacity in efficient optical networks," in *Proc. OptoElectron. Commun. Conf.*, 2011, pp. 32–33.

[15] J. Koomey, K. Kawatomo, B. Nordman, M. Peitte, and R. Brown, *Initial Comments on 'The Internet Begins With Coal'*, LBNL 44698, 1999.

[16] K. Kawatomo, J. G. Koomey, B. Nordman, R. E. Brown, M. A. Piette, M. Ting, and A. K. Meier, *Electricity Used by Office Equipment and Network Equipment in the US*, LBNL 45917, 2001.

[17] V. Turk, "Assessing the resource intensity of the internet infrastructure," M.S. thesis, IIIEE, Lund Univ., Lund, Sweden, 2001.

[18] C. Barthel, S. Lechtenbohmer, and S. Thomas, "GHG emission trends of the internet in Germany," in *Proc. IGES & Wuppental Inst. Joint Workshop Int. Climate Policy IT Sector*, pp. 55–67, 2001.

[19] K. Roth, F. Goldstein, and J. Kleinman, "Energy consumption by office and telecommunications equipment in commercial buildings," U.S. Dept. Commerce, NTIS PB2002-101438, 2002.

[20] C. Cremer, "Energy consumption of information and communication technology in Germany up to 2010," Fraunhofer ISI and CEPE Project No. 28/01, 2003.

[21] C. Lange, D. Kosiankowski, R. Hulsermann, R. Weidmann, and A. Gladisch, "Energy footprint of telecommunication networks," in *Proc. Eur. Conf. Exhibit. Opt. Commun.*, 2010, DOI: 10.1109/ECOC.2010.5621088.

[22] W. Vereecken, W. Van Heddeghem, B. Puype, D. Colle, M. Pickavet, and P. Demeester, "Optical networks: How much power do they consume and how can we optimise this?" in *Eur. Conf. Exhibit. Opt. Commun.*, 2010, DOI: 10.1109/ECOC.2010.5621575.

[23] S. Jevons, *The Coal Question—Can Britain Survive*. New York: Macmillan, 1906.

[24] H. Saunders, "The Kazzoom-Brookes postulate and neoclassical growth," *Energy J.*, vol. 13, no. 4, pp. 131–148, 1992.

[25] G. Griffa, L. Radice, C. Bianco, A. Anders, B. Zhu, D. Han, P. A. Gemma, and S. Luo, "Carbon footprint of next generation fixed networks," in *Proc. Int. Telecommun. Energy Conf.*, 2010, DOI: 10.1109/INTLEC.2010.5525646.

[26] R.-J. Essiambre, G. J. Foschini, G. Kramer, and P. J. Winzer, "Capacity limits of information transport in fiber-optic networks," *Phys. Rev. Lett.*, vol. 101, 2008, 163901.

[27] R. Tucker, "Green optical communications—Part I: Energy limitations in transport," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 261–274, Mar./Apr. 2011.

[28] M. Feng, K. Hinton, R. Ayre, and R. Tucker, "Reducing NGN energy consumption with IP/SDH/WDM," in *Proc. Int. Conf. Energy-Efficient Comput. Netw.*, 2010, DOI: 10.1145/1791314.1791344.

[29] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Architectures for energy-efficient IPTV networks," in *Proc. Conf. Opt. Fiber Commun.*, 2009, Paper OTHQ5.

[30] K. Guan, D. Kilper, and G. Atkinson, "Evaluating the energy benefit of dynamic optical bypass for content delivery," in *Proc. IEEE Conf. Comput. Commun. Workshop*, 2011, pp. 313–318.

[31] S. Figuerola, M. Lemay, V. Reijs, M. Savoie, and B. S. Arnaud, "Converged optical network infrastructures in support of future internet and grid services using IaaS to reduce GHG emissions," *J. Lightw. Technol.*, vol. 27, no. 12, pp. 1941–1946, Jun. 15, 2009.

[32] S. Lloyd, "Ultimate physical limits to computation," *Nature*, vol. 406, pp. 1047–1054, 2000.

[33] T. Li and M. C. Teich, "Bit error rate for a lightwave communication system incorporating an erbium doped fibre amplifier," *Electron. Lett.*, vol. 27, pp. 598–600, 1991.

[34] Y. Yamamoto and H. Haus, "Preparation, measurement, and information capacity of optical quantum states," *Rev. Mod. Phys.*, vol. 58, pp. 1001–1020, 1986.

[35] Alcatel-Lucent, *Wavestar OLS 40 G*, Release 3.3, 1999.

[36] Alcatel-Lucent, *Alcatel-Lucent 1830 Photonic Service Switch (PSS-64 and PSS-2 36) Data Sheet*, 2011.

[37] D. C. Kilper, R. Bach, D. J. Blumenthal, D. Einstein, T. Landolsi, L. Ostar, M. Preiss, and A. E. Willner, "Optical performance monitoring," *J. Lightw. Technol.*, vol. 22, no. 1, pp. 294–304, Jan. 2004.

[38] A. Patel, "First shared path protection scheme for generalized network connectivity in gridless optical WDM networks," in *Proc. Asian Commun. Photon. Conf. Exhibit.*, 2010, Paper PD6.

[39] P. Winzer and R.-J. Essiambre, "High-speed and high-capacity optical transmission systems," in *High Spectral Density Optical Communication Technologies*, M. Nakazawa, K. Kikuchi, and T. Miyazaki, Eds. Berlin, Germany: Springer-Verlag, 2010, pp. 102–127.

[40] M. Karlsson and E. Agrell, "Power-efficient modulation schemes," in *Impact of Nonlinearities on Fiber Optic Communications*, S. Kumar, Ed. New York: Springer-Verlag, 2010, pp. 219–253.

[41] P. Winzer, "Energy-efficient optical transport capacity scaling through spatial multiplexing," *IEEE Photon. Technol. Lett.*, vol. 23, no. 13, pp. 851–853, Jul. 1, 2011.

[42] M. Karlsson and E. Agrell, "Power-efficient modulation formats in coherent transmission systems," *J. Lightw. Technol.*, vol. 27, no. 22, pp. 5115–5126, Nov. 15, 2009.

[43] M. Karlsson and E. Agrell, "Which is the most power-efficient modulation formats in optical links?" in *Opt. Exp.*, vol. 17, pp. 10814–10819, 2009.

[44] P. Monti, P. Wiatr, A. Jirattigalachote, and L. Wosinka, "Trading power savings for blocking probability in dynamically provisioned WDM networks," in *Proc. Int. Conf. Transparent Opt. Netw.*, Munich, Germany, 2010, DOI: 10.1109/ICTON.2010.5548944.

[45] L. Chiaraviglio, D. Ciullo, M. Mellia, and M. Meo, "Modeling sleep modes gains with random graphs," in *Proc. IEEE Conf. Comput. Commun. Workshop*, 2011, pp. 355–360.

[46] D. C. Kilper and C. A. White, "Amplifier issues for physical layer network control," in *Optically Amplified WDM Networks*, J. Zyskind and A. Srivastava, Eds. Amsterdam, The Netherlands: Elsevier, 2011, pp. 221–251.

[47] D. C. Kilper, C. A. White, and S. Chandrasekhar, "Control of channel power instabilities in constant-gain amplified transparent networks using scalable mesh scheduling," *J. Lightw. Technol.*, vol. 26, no. 1, pp. 108–113, Jan. 1, 2008.

[48] S. Namiki, T. Kurosu, K. Tanizawa, J. Kurumida, T. Hasama, H. Ishikawa, T. Nakatogawa, M. Nakamura, and K. Oyamada, "Ultrahigh-definition video transmission and extremely green optical networks for future," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 446–457, Mar./Apr. 2011.

[49] P. Molinero-Fernandez and N. McKeown, "Performance of circuit switching in the internet," *J. Opt. Netw.*, vol. 2, no. 4, pp. 83–96, 2003.

[50] X. Dong, T. El-Gorashi, and J. Elmirghani, "IP over WDM networks employing renewable energy sources," *J. Lightw. Technol.*, vol. 29, no. 1, pp. 3–14, Jan. 1, 2011.

[51] G. J. Eilenberger, S. Bunse, L. Dembeck, U. Gebhard, F. Lichmann, W. Lautenschlaeger, and J. Milbrandt, "Energy-efficient transport for the future internet," *Bell Labs Tech. J.*, vol. 15, no. 2, pp. 17–168, 2010.

[52] D. C. Kilper, D. Neilson, D. Stiliadis, D. Suvakovic, and S. Goyal, "Fundamental limits on energy use in optical networks," in *Proc. Eur. Conf. Exhibit. Opt. Commun.*, 2010, DOI: 10.1109/ECOC.2010.5621584.

[53] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Green cloud computing: Balancing energy in processing, storage, and transport," *Proc. IEEE*, vol. 99, no. 1, pp. 149–167, Jan. 2011.

[54] N. I. Osman, T. El-Gorashi, and J. Elmirghani, "Reduction of energy consumption of video-on-demand services using cache size optimization," in *Proc. 8th Int. Conf. Wireless Opt. Commun. Netw.*, 2011, DOI: 10.1109/WOCN.2011.5872923.

[55] U. Lee, I. Rimac, D. Kilper, and V. Hilt, "Toward energy-efficient content dissemination," *IEEE Network*, vol. 25, no. 2, pp. 14–19, Mar./Apr. 2011.

[56] K. Guan, G. Atkinson, D. Kilper, and E. Gulsen, "On the energy efficiency of content delivery architectures," in *Proc. Int. Conf. Commun. Workshops*, 2011, DOI: 10.1109/iccw.2011.5963557.

[57] R. Kothiya, V. Tarasov, P. Sehgal, and E. Zadok, "Energy and performance evaluation of lossless file data compression on server systems," in *Proc. SYSTOR, The Israeli Exp. Syst. Conf.*, 2009, DOI: 10.1145/1534530.1534536.

[58] A. Ahmad, A. Bianco, E. Bonetto, D. Cuda, G. G. Castillo, and F. Neri, "Power-aware logical topology design heuristics in wavelength routing networks," in *Proc. 15th Int. Conf. Opt. Netw. Design Model.*, 2011, pp. 1–6.

[59] S. Azodolmolky, M. Klinkowski, E. Marin-Tordera, D. Careglio, J. Sole-Pareta,

and I. Tomkos, "A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks," *Comput. Netw.*, vol. 53, no. 7, pp. 926–944, 2009.

[60] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, "Power-aware routing and wavelength assignment in optical networks," in *Proc. Eur. Conf. Opt. Commun.*, 2009, pp. 1–2.

[61] K. Christopdoulopoulos, K. Manousakis, and E. Varvarigos, "Adapting the transmision reach in mixed line rate WDM transport networks," in *Proc. Int. Conf. Opt. Netw. Design Model.*, 2011, pp. 1–6.

[62] A. Nag, M. Tornatore, and B. Mukherjee, "Optical network design with mixed line rates and multiple modulation formats," *J. Lightw. Technol.*, vol. 28, no. 4, pp. 466–475, Feb. 15, 2010.

[63] P. Chowdhury, M. Tornatore, and B. Mukherjee, "On the energy efficiency of mixed-line-rate networks," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3.

[64] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Elastic bandwidth allocation in flexible OFDM-based optical networks," *J. Lightw. Technol.*, vol. 29, no. 9, pp. 1354–1366, May 1, 2011.

[65] A. Klekamp, F. Buchali, R. Dischler, and F. Ilchmam, "Comparison of DWDM network topologies with bitrate adaptive optical OFDM regarding restoration," in *Proc. Eur. Conf. Opt. Commun.*, 2010, pp. 1–3, DOI: 10.1109/ECOC.2010.5621398.

[66] O. Rival, A. Morea, and J.-C. Antona, "Optical network planning with rate-tunable NRZ transponders," in *Proc. Eur. Conf. Opt. Commun.*, 2009, pp. 1–2.

[67] W. Shieh, "OFDM for flexible high-speed optical networks," *J. Lightw. Technol.*, vol. 29, no. 10, pp. 1560–1577, May 15, 2011.

[68] A. Klekamp, O. Rival, A. Morea, R. Dischler, and F. Buchali, "Transparent WDM network with bitrate tunable optical OFDM transponders," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3.

[69] D. Neilson, "Photonics for switching and routing," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 4, pp. 669–678, Jul./Aug. 2006.

[70] D. J. Blumenthal, J. Barton, N. Beheshti, J. E. Bowers, E. Burmeister, L. A. Coldren, M. Dummer, G. Epps, A. Fang, Y. Ganjali, J. Garcia, B. Koch, V. Lal, E. Lively, J. Mack, M. Mašanović, N. McKeown, K. Nguyen, S. C. Nicholes, H. Park, B. Stamenic, A. Tauke-Pedretti, H. Poulsen, and M. Sysak, "Integrated photonics for low-power packet networking," *J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 458–471, Mar./Apr. 2011.

[71] S. Yoo, "Energy efficiency in the future internet: The role of optical packet switching and optical label switching," *J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 406–418, Mar./Apr. 2011.

[72] R. Tucker, "Green optical communications—Part II: Energy limitations on networks," *J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 261–271, Mar./Apr. 2011.

[73] D. Chiaroni, G. Buforn Santamaria, C. Simonneau, S. Etienne, J.-C. Antona, S. Bigo, and J. Simsarian, "Packet OADM for the next generation of ring networks," *Bell Labs Tech. J.*, vol. 14, no. 4, pp. 243–264, 2010.

[74] R. Tucker, "Optical packet switching: A reality check," *Opt. Switching Netw.*, vol. 5, pp. 2–9, 2008.

[75] K. Hinton, G. Raskutti, P. M. Farrell, and R. S. Tucker, "Switching energy and device size limits on digital photonic signal processing technologies," *IEEE J. Sel. Topics Quantum. Electron.*, vol. 14, no. 3, pp. 938–945, May/Jun. 2008.

[76] R. Tucker and K. Hinton, "Energy consumption and energy density in optical and electronic signal processing," *IEEE Photon. J.*, vol. 3, no. 5, pp. 821–833, Oct. 2011.

[77] R. Barry and P. Humblet, "Models of blocking probability in all-optical networks with and without wavelength changers," in *Proc. 14th Annu. Joint Conf. IEEE Comput. Commun. Soc.*, Boston, MA, 1995, pp. 402–412.

[78] M. Murakami, "Analyzing power consumption in optical cross-connect equipment for future large capacity optical networks," *J. Netw.*, vol. 5, no. 11, pp. 1254–1259, 2010.

[79] A. L. Chiu, G. Choudhury, G. Clapp, R. Doverspike, J. W. Gannett, J. G. Klincewicz, G. Li, R. A. Skoog, J. Strand, A. Von Lehmen, and D. Xu, "Network design and architectures for highly dynamic next-generation IP-over-optical long distance networks," *J. Lightw. Technol.*, vol. 27, no. 12, pp. 1878–1890, Jun. 15, 2009.

[80] F. Smyth, D. C. Kilper, S. Chandrasekhar, and L. P. Barry, "Applied constant gain amplification in circulating loop experiments," *J. Lightw. Technol.*, vol. 27, no. 21, pp. 4686–4697, Nov. 1, 2009.

[81] M. D. Feuer, D. C. Kilper, and S. Woodward, "ROADMs and their system applications," in *Optical Fiber Communications V B*, I. P. Kaminow, T. Li, and A. E. Willner, Eds. San Diego, CA: Academic, 2008.

[82] Cisco, *Hyperconnectivity and the Approaching Zettabyte Era*, 2009.

[83] A. J. Seeds, "Microwave-over-fiber systems," in *Optical Fiber Telecommunications VB*. London, U.K.: Elsevier, 2008, pp. 739–764.

[84] ITU-T G.707, *Interfaces for the Optical Transport Network*, 2009.

[85] M. Feng, R. Ayre, K. Hinton, and R. Tucker, "Energy consumption in intelligent optical networks," in *GreenIT*, DOI: 10.5176/978-981-08-7240-3_G-22, 2010.

ABOUT THE AUTHORS

Daniel Kilper (Senior Member, IEEE), received the B.S. degree in electrical engineering and the B.S. degree in physics with honors from Virginia Polytechnic Institute & State University, Blacksburg, in 1990 and the M.S. and Ph.D. degrees in physics from The University of Michigan, Ann Arbor, in 1992 and 1996, respectively.

Following his Ph.D., he was a Postdoctoral Research Scientist at the Montana State University Optical Technology Center in the Department of Physics. He became an Assistant Professor of Physics at the University of North Carolina, Charlotte, in 1997. In 2000, he joined the Advanced Photonics Research Department as a Member of Technical Staff at Bell Labs, Lucent Technologies, Holmdel, NJ, now Alcatel-Lucent. He holds six patents and authored three book chapters and 91 peer-reviewed journal articles and published proceedings. While at Bell Labs he has conducted research on optical performance monitoring and on transmission, architectures and control systems for transparent and energy-efficient optical networks.

Dr. Kilper serves as the chair of the Technical Committee in the GreenTouch Consortium and is the Bell Labs Liaison Executive on the Operations Committee of the Center for Energy Efficient Telecommunications, Melbourne, Australia. He is an associate editor for the OSA/IEEE JOURNAL OF OPTICAL COMMUNICATIONS AND NETWORKING. He served on technical program committees for ICTON, INFOCOM GCN Workshop, IFIP Networking, IQEC, CLEO/Europe, COIN/ACOFT, and Optical Technology in the Carolinas Conference and was General Chair of the OPTO-Southeast Conference. He has been an invited panelist, presenter, and workshop organizer at OFC/NFOEC. During 2003-2006 he was a member of the Bell Labs Advisory Council on Research. He received the Bell Laboratories President's Gold Medal Award in 2004. He is a member of the Optical Society of America. He has served on the OSA Member and Educational Services Committee, Beller Award Committee, and chaired the OSA Leadership Award Committee.

K. Guan received the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 2007.

From 2007 to 2010, he was a Senior Research Engineer at BAE Systems, Wayne, NJ. He joined Bell Labs, Alcatel-Lucent, NJ, in 2010. His current research interests include energy-efficient network architectures, optical transmission systems, and network optimization.

K. Hinton, photograph and biography not available at the time of publication.

R. Ayre, photograph and biography not available at the time of publication.